Shaping the Future of Mathematics with AI
AIM 2027 Vision for the Netherlands

Christoph Brune
AIM chair

Tim van Erven
AIM co-chair

Version 1.0, November 27, 2026, www.ai4math.nl

Introduction

2026 will likely be remembered as the year in which mathematics changed. Several long-standing problems, some of which had resisted generations of effort, have been resolved with substantial help from AI systems, and research-level proofs are now produced at a pace and cost that would have seemed implausible only a few years ago. Prominent examples are the Erdős unit distance conjecture[1], the forced[2] and unforced[3] Navier–Stokes problems, and most recently the Riemann hypothesis[4].

The question is no longer whether AI will be part of the future of mathematics, but what kind of mathematics we want to build with it. It is tempting to respond either with uncritical enthusiasm or with defensive alarm. The vision of AIM[5] is more nuanced, addressing points that need urgent action while recognizing the opportunities that AI offers.

Terence Tao has given this moment a useful name. For most of its history, mathematics operated under *proof scarcity*: correct proofs of hard results were rare and costly, and our journals, careers, prizes and PhD programs grew up around rewarding them. With AI we are entering an era of *proof abundance*, and Tao calls for a deliberate move towards 'Math 2.0', in which AI strengthens human understanding rather than bypassing it.[6] Solving problems, in this view, was always a proxy for deeper goals: insight, new concepts, and maps of the mathematical landscape that others can use. With proofs in abundance, these goals move to the center: digesting and explaining results, exploring what lies around a theorem, building shared formal libraries, and collaborating at previously unimaginable scales. The main bottleneck is no longer finding proofs, but our imagination.

Whatever role AI comes to play in generating arguments or checking details, the responsibility for understanding a result, judging its value, and standing behind its correctness has to *remain human*. This keeps mathematics valuable to other disciplines and to society, and keeps it something students and researchers want to do. The joy of discovery, the individual voice and taste behind a research program, and a community that thinks and argues together are not a bonus: they are why our discipline is worth teaching and why its results are convincing and trustworthy. Wherever mathematics informs consequential decisions, in industry, policy or elsewhere, the result cannot simply be an AI's output; it needs a human who understands it well enough to explain and stand behind it.

For students and early-career mathematicians, machines that can already reason like a professional mathematician raise the question of what their own role will be. We believe it can be a rich one,

and creating the room for it is the responsibility of those of us who shape the field through hiring, curricula, editorial boards and funding. For the first time in generations, the habits and institutions of mathematics are open for redesign: which questions are worth asking, how results are digested and shared, and how trust and the joy of mathematics are preserved. This requires the experience of *established mathematicians* as much as the *fresh perspective* of those entering the field, who deserve a real voice in these decisions. Together, across generations, we can shape this future ourselves.

Building on the Leiden Declaration on Artificial Intelligence and Mathematics, this document is a starting point for discussion within the Dutch mathematical community, moving from what remains uniquely human, via the new horizons AI opens up, to the changes needed in training, research practice, publishing and infrastructure.

1 Affirm the humanity of mathematics

Value, trust, teaching and community in mathematics are uniquely human.

1.1 Value comes from the effect on people

A mathematical result acquires value through its effect on the mathematical community and on society. AI working in a vacuum cannot produce such value by itself.

> For example, the recent AI-assisted solutions to multiple Millennium Prize Problems are exciting not because the problems were difficult — there exist many difficult but uninteresting problems — but because of a respected community of people who care deeply about these particular mathematical questions, and who understand their significance in light of a long cultural tradition.

Recommendations

1. Human mathematicians must curate the canon of mathematics: what is 'interesting' and why?
2. Mathematics is infinite, and even with AI we cannot explore all of it. Humans must determine at which mathematical questions to direct our limited resources.
3. Digesting and communicating (AI-assisted) mathematical results will become much more important. We should recognize this early, and give such work higher status.
4. Humans must decide which concepts and techniques to teach our students.

1.2 Trust and accountability rely on people

Only human mathematicians can hold trust and be accountable to society when choosing mathematical models and trade-offs in high-stakes applications.

> Mathematics makes things reliable. For instance, when securing communication with cryptography or testing vaccine effectiveness with statistics. But this often involves modeling choices, trade-offs between conflicting interests, and pragmatic leaps from mathematical theory to practice. These steps necessarily require human experts, because we understand when and how we can trust humans, and we have societal mechanisms to hold them accountable. None of these mechanisms exist for AI, which therefore cannot hold responsibility.

Recommendations

1. Modeling choices and trade-offs between conflicting interests and objectives must be made by humans.
2. We need human experts who deeply understand the math underlying high-stakes applications; otherwise they cannot take responsibility, or communicate effectively.

1.3 Teaching requires human contact

Students need personal contact with teachers for guidance and inspiration, and connections with fellow students for well-being.

> If designed well, personalized AI tutors hold great promise for learning. But we have seen during the Covid pandemic how students suffered from studying remotely, without direct contact with fellow students or with teachers to motivate and inspire them. This shows that learning is not just a technical activity, but crucially relies on community and human interaction.

Recommendations

1. Experiment with AI-supported learning, but do not let it replace human contact with teachers.
2. Make regular student interaction with other students an explicit goal in teaching programs.

1.4 Human agency in mathematics rests on human communities

Safeguarding human agency in the context of AI relies crucially on people and the vitality of the mathematical communities that support them.

> There are many mathematical subcommunities that are vital to the health of mathematics: teachers, students, researchers, entrepreneurs, hobbyists, etc. They are all needed to shape the future of mathematics. Communication, personal relationships, and feeling connected are especially important in times of change.

Recommendations

1. Organize community events to discuss and digest changes caused by AI.
2. Ensure broad representation, including from groups that are often underrepresented.
3. When making (policy) changes, consider how they will change incentives and how this will impact various groups and subcommunities. Pay particular attention to the effect on underrepresented groups.

2 New horizons for mathematics with AI

As AI changes the way we do mathematics, exciting new opportunities emerge for teachers, mathematical engineers, and researchers.

2.1 Teachers can be empowered by AI

AI already makes it easy to create interactive animations and custom presentation materials. There is also much potential to use it to automate routine parts of grading, and to develop personalized AI-assisted tutoring. When used with good judgment, AI can make teaching better for students, and the teacher's job more enjoyable, allowing them to spend more time with students.

Recommendations

1. Take advantage of AI tools to easily develop customized teaching materials that are more effective and more engaging.
2. Pioneer AI tutors that support learning, e.g. via custom AI personas that draw on all class reading materials but nudge students to address gaps in their understanding instead of giving away answers too easily.
3. Develop reliable ways to use AI to support straightforward parts of grading, e.g. for grading exercises that do not require human judgment to evaluate.

2.2 Mathematical engineers can be empowered by AI

Mathematicians in industry can apply a broader set of mathematical techniques with AI, and tap into new research more easily.

Recommendations

1. Train students and retrain working mathematical engineers to develop productive AI skills that empower them. See also point 3.2.

2.3 Mathematical researchers can be empowered by AI

In the short term, AI disrupts many practices of mathematical research. But in the long term, it can open up a wealth of new mathematical applications, allow us to explore new frontiers of understanding, create opportunities for young people to reshape how we work, and make the field more self-aware about the many different reasons for doing mathematics.

Recommendations

1. Technical work that can easily be done by AI with a single prompt will not remain publishable by itself. It will still be worthwhile to do such work by hand to build up deep domain understanding and for enjoyment, but not as final scientific output. Human contributions remain of great value, for instance in theory building and the conceptual part of research, or in technical questions (old or new) that are hard even with AI support.
2. AI will not make mathematics run out of hard problems: new and exciting questions will come up that cannot be solved even with AI support. Seize opportunities to identify and explore such questions.
3. In some parts of applied mathematics, developing mathematical theory gives a deep understanding, but is much slower than trying out ideas in experiments. For instance, in machine learning, this has led to a widening gap between mathematical theory and the techniques used in practice. Exploit faster theory development with AI to close such gaps.
4. With AI tools, new mathematical theory can be turned into proof-of-concept demonstrations with a high technology readiness level much more quickly. Take advantage of such tools to open up a wide range of new mathematical applications, with corresponding funding opportunities.
5. We are going through a time of change: many more people than ever are open to changing old processes and habits. Invite young researchers with fresh, out-of-the-box ideas to help shape the future of how we work.
6. As argued by Tao,[5] AI developments force us to separate the many reasons for doing mathematics. Make these reasons explicit: this is worth doing regardless of the role of AI. Russell's paradox and Gödel's incompleteness theorems similarly forced mathematicians to reexamine foundational assumptions that had previously been left implicit, and mathematics emerged stronger for it.

3 Skills, training and evaluation

Update training and evaluation to ensure that AI amplifies human mathematical ability.

3.1 Deep mathematical insight remains crucial

Deep understanding of a mathematical topic is the foundation of being a mathematician. But we can no longer judge understanding from written work (papers, PhD theses) alone. Nor does producing such work guarantee that the author builds up deep insight.

Recommendations

1. Evaluate mathematical insight more thoroughly when hiring and throughout the PhD process, for instance by an in-depth discussion about an applicant's recent work.

3.2 AI requires new skills

Using AI effectively requires many additional skills. This extends beyond technical proficiency to, for instance, staying in control and adopting effective workflows that include human learning.

Recommendations

1. Include both technical AI skills and self-reflection on AI use as mandatory components in the teaching curriculum. This includes, for instance, preventing cognitive offloading of valuable learning opportunities to AI, and exercising human judgment in the presence of AI.
2. Organize seminars about AI skills for mathematicians, mathematical engineers and mathematics teachers in all sectors.

3.3 Revalue communication and digestion

AI use will rapidly increase the amount of 'dark math': mathematical results that are proven somewhere, but that nobody understands or knows about.[7] Extracting important insights and communicating them will therefore become increasingly important.

Recommendations

1. Place higher value on digestion and communication of mathematical results in hiring decisions.
2. Extend the concept of mathematical discovery to digesting dark math.

3.4 PhD programs need to be significantly updated

The goal of a PhD is to become an independent mathematical researcher. In the past, we have evaluated this mostly based on the ability to write a good PhD thesis, but this is no longer sufficient. A PhD has always been about a broader set of skills, but we now need to articulate these skills more explicitly, and significantly update our processes to ensure PhD students acquire them.

Recommendations

Build on existing proposals, in particular the report of the Harvard summit on PhD math education in the age of AI.[8] Three of its core ideas translate directly into recommendations for Dutch PhD programs:

1. Make the *broader skill set* explicit in the PhD. Asking good questions, collaborating, working across disciplines, and using computational tools, AI, and formalization wisely should be trained deliberately. Communicating mathematics to experts and broad audiences should become an explicit degree expectation supported for example by group meetings and student seminars.
2. *Assess the person*, not only the text. A PhD should not be awarded solely based on the written dissertation. For example, oral assessments by several faculty members can provide meaningful feedback throughout the PhD. The process should end in a rigorous defense showing that the candidate fully masters, and truly owns, the thesis results.
3. Use AI to *accelerate human understanding*. Advisors and students should discuss AI use regularly and openly. Three principles apply: authors remain fully responsible for the correctness and understanding of everything under their name, AI use is always disclosed, and material is only shared with an AI system with everyone's consent. Students should still learn to work by hand and think long about hard problems.

4 Overhaul publishing and reviewing

Experiment boldly with reforms of paper and grant reviewing.

4.1 The current paper review system is untenable

AI makes writing a paper much easier. This is causing a massive increase in the number of submissions. Existing review systems must adapt to avoid getting overloaded.

Recommendations

1. Experiment with new reviewing systems and output formats to discover what works and what does not.
2. Learn from reviewing experiments in other fields (for instance in machine learning and theoretical computer science, where increasing numbers of submissions have been a problem for a longer time). See, for instance, STOC 2027[9], which combines several experimental features.
3. Keep writing results up in human-readable papers or other human-accessible forms: results that are accessible only through AI mediation lock us into using a rapidly changing technology with an unpredictable future.
4. Insist on high standards of rigor and communication. Standards can go up if AI support is used appropriately, for example via automatic Lean formalizations.
5. Require human authors, who take responsibility for paper quality, significance, and attribution.
6. Require AI use statements to make the role of AI explicitly visible.
7. Use the current momentum for reform to switch to open access publishing.

4.2 The current grant review system is untenable

AI makes writing a grant proposal much easier. This is causing a massive increase in the number of submissions. Existing review systems must adapt to avoid getting overloaded.

Recommendations

1. Experiment more with new reviewing systems to discover what works and what does not.
2. Recognize and stimulate the expertise already available, for instance at NWO, about alternative reviewing procedures.

5 Invest in AI model choice and infrastructure to support it

Different tasks require different AI models. Develop national infrastructure to support model diversity, for security, privacy, and to reduce dependence on a few non-European providers.

5.1 Know the right model for each task

Different tasks in research, teaching and project management have different requirements for model abilities, data privacy, computer security, energy usage, costs, etc. Develop national expertise in model choice.

Recommendations

1. Use different AI models for different tasks: simple tasks might be handled by efficient models that require less energy; privacy-sensitive tasks might need strict data control or European models like Mistral; frontier proofs might, for now, require American commercial models.
2. Develop national expertise in model choice, for instance via benchmarks that measure model abilities

for the actual tasks mathematicians encounter in practice. These might include, for instance, developing new proofs and counterexamples, automatic proof formalization in Lean, transforming Lean proofs into human-understandable proofs, programming, creating figures and presentations, checking informal proofs, mathematical writing, or project management.

5.2 Develop nationally supported infrastructure that provides secure access to all required models, including agentic AI

This should cover the full spectrum from energy-efficient models and open-weight models (Chinese and European) to state-of-the-art commercial models (from the US). For many mathematical tasks, agentic AI is essential.

Recommendations

1. Build on existing expertise at, for instance, SURF AI and UvA AI Chat. This includes technical expertise, expertise in data management best practices, and additional customizations like AI personas for teaching.
2. Insist on including access to agentic AI with top commercial models, which is currently necessary for the most challenging mathematical proofs.
3. Ensure stable long-term funding. The costs of running open-weight models, e.g. at SURF, will go up with more usage, and prices for commercial access may go up when tech companies stop sponsoring model access as part of their growth strategy.
4. Organize access for all scientific disciplines together at the national level with all universities.
5. Agentic AI is important, but running it locally without proper precautions can cause significant IT security risks. Organize centralized development of IT best practices (e.g. using sandboxes or dedicated computers) for agentic AI at a national level.

Acknowledgments

We thank Jim Portegies for constructive critical comments from the perspective of the Leiden Declaration.

Notes

1. OpenAI, *Planar Point Sets with Many Unit Distances*, May 2026, https://cdn.openai.com/pdf/74c24085-19b0-4534-9c90-465b8e29ad73/unit-distance-proof.pdf. N. Alon et al., *Remarks on the Disproof of the Unit Distance Conjecture*, arXiv:2605.20695, 2026, https://arxiv.org/abs/2605.20695.

2. OpenAI, *On the Navier–Stokes Millennium Prize Problem*, Sep 2026, https://openai.com/index/navier-stokes-solution/. *Quanta Magazine*, AI Has Solved One of Math's $1 Million Millennium Prize Problems, 8 Sep 2026, https://www.quantamagazine.org/ai-has-solved-one-of-maths-1-million-millennium-prize-problems-20260908/.

3. L. Alpöge, T. Buckmaster, *Finite-Time Blow-Up for the Navier–Stokes Equations*, 20 Nov 2026, with Lean 4 formalization, https://github.com/tristanbuckmaster/navier-stokes/. *Scientific American*, AI Just Solved the 'Real' Navier–Stokes Problem, 2026, https://www.scientificamerican.com/article/ai-just-solved-the-real-navier-stokes-problem/.

4. L. Alpöge, *Lean Proof of the Riemann Hypothesis and the Generalized Riemann Hypothesis for Dirichlet L-functions*, 20 Nov 2026, https://alpo.ge. Lean FRO and Mathlib Initiative, *Lean Verification of the Riemann Hypothesis*, 2026, https://leanprover-community.github.io/blog/posts/riemann-hypothesis/.

5. Dutch AI & Mathematics initiative, https://www.aimath.nl.

6. T. Tao, Can maths survive AI? Fields medallist Terence Tao calls for a 'Maths 2.0' world, *Nature*, World View, Oct 2026; see also his Math 2.0 lecture (Caltech, Oct 2026) and Mathematics in the age of AI.

7. By analogy with 'dark matter' in astronomy.

8. Summit on PhD Math Education in the Age of AI, Harvard, September 17–18, 2026, https://cmsa.fas.harvard.edu/aimathphd_summit/.

9. STOC 2027 call for papers, https://acm-stoc.org/stoc2027/stoc2027-cfp.html.